arXiv:1610.06546v1 [quant-ph] 20 Oct 2016

Hamiltonian Simulation by Qubitization

Guang Hao Low[*], Isaac L. Chuang[†]

October 21, 2016

Abstract

Given a Hermitian operator $\hat{H} = \langle G|\hat{U}|G\rangle$ that is the projection of an oracle $\hat{U}$ by state $|G\rangle$ created with oracle $\hat{G}$, the problem of Hamiltonian simulation is approximating the time evolution operator $e^{-i\hat{H}t}$ at time t with error ϵ. We show that this can be done with query complexity $\mathcal{O}(t + \log(1/\epsilon))$ to $\hat{G}, \hat{U}$ that is optimal not just in asymptotic limits, but for all values t, ϵ. Furthermore, only 2 additional ancilla qubits are required in total, together with $\mathcal{O}(1)$ additional single and two-qubit gates per query. Our approach to Hamiltonian simulation subsumes important prior art considering Hamiltonians which are d-sparse or a linear combination of unitaries, leading to significant improvements in space complexity, as well as a quadratic speed-up for precision simulations. It also motivates useful new instances, such as where $\langle G|\hat{U}|G\rangle$ is a density matrix. A key technical result is 'qubitization' which uses controlled-$\hat{U}$ and controlled-$\hat{G}$ to embed $\hat{H}$ in an invariant SU(2) subspace. A large class of operator functions of $\hat{H}$ can then be computed with optimal query complexity, of which $e^{-i\hat{H}t}$ is a special case.

1 Introduction

Quantum computers were originally envisioned as machines for efficiently simulating quantum Hamiltonian dynamics. As Hamiltonian simulation is BQP-complete, the problem is believed to be intractable by classical computers, and remains a strong primary motivation. The first explicit quantum algorithms were discovered by Lloyd [31] for local interactions, and then generalized by Aharonov and Ta-Shma [3] to sparse Hamiltonians. Celebrated achievements over the years [20, 8, 32, 10, 33] have each ignited a flurry of activity in diverse applications from quantum algorithms [26, 19, 21, 15] to quantum chemistry [46, 44, 40, 41, 4, 29, 39]. In this dawning era of the small quantum computer [5, 23], the relevance and necessity of space and gate efficient procedures for practical Hamiltonian simulation has intensified.

The cost of simulating the time evolution operator $e^{-i\hat{H}t}$ depends on several factors: the number of system qubits n, evolution time t, target error ϵ, and how information on the Hamiltonian $\hat{H}$ is made available. This field has progressed rapidly following groundbreaking work in the fractional query model [8] achieving query complexities that depend logarithmically on error. This was generalized by Berry, Childs, Cleve, Kothari, and Somma (BCCKS) [9] to the case where $\hat{H} = \sum_{j=1}^{d} \alpha_j \hat{U}_j$ is a linear combination of d unitaries and the $|\alpha_j|$ sum to α – such a decomposition always exists – with an algorithm using $\mathcal{O}(\frac{\log(d)\log(\alpha t/\epsilon)}{\log\log(\alpha t/\epsilon)})$ ancilla qubits and only $\mathcal{O}(\frac{\alpha t \log(\alpha t/\epsilon)}{\log\log(\alpha t/\epsilon)})$ queries. Subsequently [10], an extension to d-sparse Hamiltonians was made, where $\hat{H}$ has $\leq d$

[*]Department of Physics, Massachusetts Institute of Technology {glow@mit.edu}

[†]Department of Electrical Engineering and Computer Science, Department of Physics, Research Laboratory of Electronics, Massachusetts Institute of Technology. {ichuang@mit.edu}

Algorithm	Model	Ancilla qubits	Query Complexity $\mathcal{O}(\cdot)$	Gates per Query $\mathcal{O}(\cdot)$	
				Oracle	Primitive
Sparse [33]	$\hat{H}_{jk}$	$n+m+3$	$dt\|\hat{H}\|_{\max}+\log(1/\epsilon)$	Varies	$n+m\text{polylog}(m)$
BCCKS [9]	$\sum_j \alpha_j \hat{U}_j$	$\mathcal{O}(\frac{\log(d)\log(\alpha t/\epsilon)}{\log\log(\alpha t/\epsilon)})$	$\frac{\alpha t\log(\alpha t/\epsilon)}{\log\log(\alpha t/\epsilon)}$	dC	1
LMR [31]	Mixed ρ	$n+1$	t^2/ϵ	Varies	$\log n$
Thm. 1	$\langle G\|\hat{U}\|G\rangle$	$\lceil\log_2(d)\rceil+2$	$t+\log(1/\epsilon)$	Varies	1
Cor. 8	$\sum_j \alpha_j \hat{U}_j$	$\lceil\log_2(d)\rceil+2$	$\alpha t+\log(1/\epsilon)$	dC	1
Cor. 9	Purified ρ	$n+\lceil\log_2(d)\rceil+2$	$t+\log(1/\epsilon)$	Varies	$\log n$

Table 1: Performance comparison of state-of-art with our new approaches (bottom three lines), for Hamiltonian simulation $e^{-i\hat{H}t}$ of $\hat{H}\in\mathbb{C}^{2^n\times 2^n}$ with error ϵ. The d-sparse simulation oracle describes entries of $\hat{H}$ with maximum absolute value $\|\hat{H}\|_{\max}$ to m bits of precision. The BCCKS oracle provides the decomposition $\hat{H}=\sum_{j=1}^{d}\alpha_j\hat{U}_j$, $\|\hat{H}\|\leq\alpha=\sum_j|\alpha_j|$ and each $\hat{U}_j$ has cost $\mathcal{O}(C)$. The LMR query complexity refers to samples of the density matrix $\hat{\rho}=\hat{H}$. This work generalizes the above with oracles $\hat{G}|0\rangle=|G\rangle\in\mathbb{C}^d,\hat{U}$ such that $\langle G|\hat{U}|G\rangle=\hat{H}$, where $\|\hat{H}\|\leq 1$. A new model where the oracle that outputs the purification $|\rho\rangle=\sum_{j=1}^{d}\alpha_j|j\rangle_a|\psi_j\rangle$, $\text{Tr}_a[|\rho\rangle\langle\rho|]=\hat{\rho}$ is provided.

non-zero elements per row with max-norm $\|\hat{H}\|_{\max}$, to achieve a quadratic improvement in sparsity with $\mathcal{O}(\frac{dt\|\hat{H}_{\max}\|\log(dt\|\hat{H}_{\max}\|/\epsilon)}{\log\log(dt\|\hat{H}_{\max}\|/\epsilon)})$ queries. A prominent open question featured in all these works was whether the *additive* lower bound $\Omega(t+\frac{\log(1/\epsilon)}{\log\log(1/\epsilon)})$ was achievable for any of these models.

Recently [33], we found for the d-sparse model, which is popular in algorithm design by quantum walks [17], a procedure achieving the optimal trade-off between all parameters, with query complexity $\Theta(dt\|\hat{H}\|_{\max}+\gamma)$, where $1\geq\frac{\gamma}{\log(1/\epsilon)}\geq\frac{1}{\log\log(1/\epsilon)}$ depends on the value of t. The strictly linear-time performance with additive complexity is a quadratic improvement for precision simulations $t\sim\log(1/\epsilon)$, and the constant number of $n+m+3$ ancilla qubits significantly improves on prior art which depends on t,ϵ like $\mathcal{O}(\log(t/\epsilon))$. The approach, based on Childs' [16, 7, 30] extension of Szegedy's quantum walk [43], required two quantum oracles: one accepting a j row and k column index to return the value of entry $\hat{H}_{jk}$ to m bits of precision, and another accepting accepting a j row and l sparsity index to return in-place the l^{th} non-zero entry in row j.

Unfortunately, the d-sparse model is less appealing in practical implementations for several reasons. First, it is exponentially slower than BCCKS when the $\hat{U}_j$ are of high weight with sparsity $\mathcal{O}(2^n)$. Second, its black-box oracles can be challenging to realize. Avoiding the $\mathcal{O}(2^n)$ blowup by exploiting sparsity requires that positions of non-zero elements are efficiently row-computable, which is not always the case. Third, the Childs quantum walk requires a doubling of the n system qubits, which is not required by BCCKS. It was unclear whether our methodology could be applied to other formulations of Hamiltonian simulation, in contrast to alternatives that seem more general [12].

Ideally, the best features of these two algorithms could be combined (Table. 1). For example, given the decomposition $\hat{H}=\sum_{j=1}^{d}\alpha_j\hat{U}_j$, one would like the optimal additive complexity of sparse Hamiltonian simulation, but with the BCCKS oracles that are more straightforward to implement. Furthermore, one could wish for a constant ancilla overhead, of say $\lceil\log_2(d)\rceil+2$, superior to either algorithm. These improvements would greatly enhance the potential of early practical applications of quantum computation.

We achieve precisely this optimistic fusion via an extremely general procedure, made possible by what we call 'qubitization', that subsumes both and motivates new formulations of Hamiltonian simulation, as captured by this main theorem:

Theorem 1 (Optimal Hamiltonian simulation by Qubitization)**.** *Given a unitary $\hat{U}$ on q qubits and state preparation unitary $\hat{G}|0\rangle = |G\rangle$ such that $\langle G|\hat{U}|G\rangle = \hat{C}$ is a complex matrix on $n \leq q$ qubits, $e^{-i\hat{H}t}$ can be simulated for time t, error ϵ in trace distance, and failure probability $\mathcal{O}(\epsilon)$, with $q+2$ qubits in total and $\Theta(N)$ queries to controlled-$\hat{G}$, controlled-$\hat{U}$, and their inverses, and $\mathcal{O}(N)$ single and two-qubit primitive quantum gates where*

$$\hat{H} = \frac{1}{2}(\hat{C} + \hat{C}^\dagger), \quad N = t + \gamma(t,\epsilon), \quad \gamma(t,\epsilon) = \begin{cases} \mathcal{O}(\log(1/\epsilon)/\log\log(1/\epsilon)), & t = \mathcal{O}(\gamma(t,\epsilon)), \\ \mathcal{O}(\log(1/\epsilon)), & \textit{otherwise.} \end{cases} \tag{1}$$

The optimality of the procedure follows by using the Childs quantum walk for $\hat{G}, \hat{U}$. Furthermore, the transparent nature of Thm. 1 significantly expedites the development of new useful formulations of Hamiltonian simulation. For instance, we easily obtain a new result for the scenario where $\hat{H}$ is a density matrix $\hat{\rho}$. Whereas $\hat{\rho}$ can be produced by discarding the ancilla of some output from a quantum circuit $\hat{G}$, we instead keep this ancilla, leading to an unconditional quadratic improvement in time scaling, and an exponential improvement in error scaling over the sample-based Lloyd, Mohseni, and Rebentrost (LMR) model [32, 28], as summarized in Table. 1. This scenario could enhance recent advances, such as quantum semidefinite programming [15].

In fact, Hamiltonian simulation is an application of our main innovation: an approach we call the 'quantum signal processor', where the equation $\langle G|\hat{U}|G\rangle = \hat{C}$ in Thm. 1 is interpreted as a non-unitary *signal* operator encoded in a subspace of an oracle $\hat{U}$ flagged by $|G\rangle$. Many problems, highlighted in Table. 2, can be of this form. Qubitization is the essential first step, where the oracles $\hat{G}, \hat{U}$ are queried to obtain a Grover-like search parallelized over each eigenvalues λ of $\hat{H} = \frac{1}{2}(\hat{C} + \hat{C}^\dagger)$. Each *iterate*, $e^{-i\hat{\sigma}_y \otimes \cos^{-1}[\hat{H}]}$, of the search is implementable deterministically with $\mathcal{O}(1)$ queries, and resembles Szegedy's [43] and Childs' [16] quantum walk. However, the key difference lies in the extremely general encoding of the signal through any $\hat{G}$ and $\hat{U}$, instead of via oracles of the d-sparse formulation.

Unlike the Grover search, we do not seek to prepare some target state. Rather, we exploit Grover-like rotations, which are isomorphic to SU(2), to engineer arbitrary target functions $f(\lambda)$ of its overlap λ. The quantum signal processor exploits this structure to attack the often-considered problem of designing a quantum circuit $\hat{Q}$ that queries $\hat{G}, \hat{U}$ such that $\langle G|\hat{Q}|G\rangle = f[\hat{H}]$ for some target function $f[\cdot]$. Whereas prior art based on the linear-combination of unitaries algorithm [30] required a detailed analysis for f on a case-by-case basis and post-selection leading to suboptimal scaling, the quantum signal processor computes $f(\hat{H})$ with an optimal query complexity that exactly matches polynomial lower bounds for a large class of functions. We provide two methods: 'observable transformations' where $\Re f(\lambda)$, $\Im f(\lambda)$ are of equal parity, and 'quantum signal processing'[1] where $\Re f(\lambda)$, $\Im f(\lambda)$ are of opposite parity. Thus generic improvements to all applications in Table. 2 can be expected in query complexity, ancilla overhead, and scope of possible signal inputs. In particular, Thm. 1 follows directly from the query complexity of the choice $f(\lambda) = e^{-i\lambda t}$, which corresponds to applying $t\sin(\cdot)$ on eigenphases of the iterate.

The implications of the quantum signal processor stretch beyond the applications of Table. 2. As Hermitian operators are also quantum observables, our methodology also offers a new approach for implementing arbitrary functions of measurement operators on quantum states [1, 13, 24, 38]. We conjecture that this application would be optimal, which is supported by our optimal results for Hamiltonian simulation. In Sec. 2, we introduce the framework for the quantum signal processor, and provide an overview of the major problems it addresses. These are then developed in detail: qubitization in Sec. 3, implementing arbitrary functions of $\hat{H}$ in Sec. 4, and finally application to Hamiltonian simulation in Sec. 5.

[1] Quantum signal processing was presented in [33] for the Childs quantum walk, and included for completeness.

Problem	BCCKS [10]	d-sparse [33]	Evolution by ρ	QPE	QLSP [19]	Gibbs [21]
$\hat{H}$	Hamiltonian	Hamiltonian	Density matrix	Unitary	Matrix	Hamiltonian
$\hat{U}$	Selects $\hat{U}_j$	Isometry $\hat{T}$	SWAP	Any	Any	Any
$\hat{G}$	$\hat{U}_j$ coefficients	Identity	Purified ρ	Any	Any	Any
Solution	$e^{-i\hat{H}t}$	$e^{-i\hat{H}t}$	$e^{-i\hat{\rho}t}$	Decision	$\hat{H}^{-1}$	$e^{-\beta\hat{H}}$

Table 2: List of six example problems (top row), solvable using the quantum signal processor approach to compute an operator function $f[\cdot]$ of $\hat{H}$, the Hermitian component of $\hat{C} = \langle G|\hat{U}|G\rangle$. Through qubitization, the scope of inputs to the Quantum Linear Systems Problem (QLSP) and Gibbs Sampling (Gibbs) can be any $\hat{H}$ of this form, either indirectly through Hamiltonian simulation, or directly through quantum signal processing. Quantum Phase Estimation (QPE) here decides whether eigenphases θ of an implemented unitary satisfy some property e.g. $f(\theta) \geq 1/2$.

2 Overview of the Quantum Signal Processor

Since coherent quantum computation is restricted to unitary operations, one commonly finds a situation such as in Table. 2, where only a portion of a given unitary *signal oracle* $\hat{U} : \mathcal{H}_s \otimes \mathcal{H}_a \to \mathcal{H}_s \otimes \mathcal{H}_a$ implements the desired dynamics of some non-unitary *signal operator* $\hat{H} : \mathcal{H}_s \to \mathcal{H}_s$ on an arbitrary n qubit input system quantum state $|\psi\rangle \in \mathcal{H}_s$. Specifically, $\hat{H} = \langle G|\hat{U}|G\rangle$ in a subspace flagged by an ancilla *signal state* $|G\rangle = \hat{G}|0\rangle \in \mathcal{H}_a$ of dimension d prepared by a unitary *state* oracle $\hat{G} \in \mathcal{H}_a$ from some standard basis. This naturally divides $\hat{U}$ into two subspaces – $\mathcal{H}_G = \text{span}\{|G\rangle \otimes \mathcal{H}_s\}$ where the desired evolution $\hat{U}|G\rangle|\psi\rangle = |G\rangle\hat{H}|\psi\rangle$ is implemented, and one for everything else. Thus

$$\hat{U}|G\rangle_a|\psi\rangle_s = |G\rangle_a\hat{H}|\psi\rangle_s + g[\hat{H}]|G_\psi^\perp\rangle_{as}, \quad g(\lambda) = \sqrt{1-|\lambda|^2}, \quad \hat{U} = \begin{pmatrix} \hat{H} & \cdot \\ \cdot & \cdot \end{pmatrix}. \tag{2}$$

where the signal operator $\hat{H}$ is in general a complex matrix with spectral norm $\|\hat{H}\| = \|\langle G|\hat{U}|G\rangle\| \leq 1$. We represent $\hat{U}$ such that the top-left block is precisely $\hat{H}$, and acts on an input state $|G_\psi\rangle \equiv |G\rangle|\psi\rangle \in \mathcal{H}_G$. Also important is the orthogonal output component in $\mathcal{H}_{G^\perp} = \text{span}\{|G_\psi^\perp\rangle, \forall|\psi\rangle \in \mathcal{H}_s\}$. In the following, we consider the case where $\hat{H}$ is a normal matrix, until otherwise stated. Furthermore, all operator functions are indicated by $[\cdot]$ where applicable and defined by $f[\hat{H}] = \sum_\lambda f(\lambda e^{i\theta_\lambda})|\lambda\rangle\langle\lambda|$ for any scalar function $f(\cdot)$, where $\hat{H}|\lambda\rangle = \lambda e^{i\theta_\lambda}|\lambda\rangle$ are eigenstates and $\lambda \in \mathbb{R}$. Whenever the context is clear, we drop the ancilla and system subscripts. We summarize these assumptions and resources, as well as other requirements as follows:

Quantum Signal Processor:

Assumptions: (1) The signal oracle $\hat{U}$ is an m qubit unitary; (2) the state oracle $\hat{G}$ is a dimension d unitary; (3) the signal state $|G\rangle = \hat{G}|0\rangle$ is prepared by $\hat{G}$ from some standard computational basis; (4) the signal operator $\hat{H} = \langle G|\hat{U}|G\rangle$ is an n qubit normal operator.

Resources : (1) Signal oracle $\hat{U}, \hat{U}^\dagger$, and their controlled versions; (2) state oracle $\hat{G}, \hat{G}^\dagger$, and their controlled versions; (3) comptational basis state $|0\rangle$; (4) $m+2$ qubits; (4) arbitrary single and two-qubit gates.

In the simplest case, one might wish to apply $\hat{H}$ multiple times to generate higher moments. When $\hat{H}$ is proportional to a unitary, this corresponds to phase accumulation, which is essential to precision measurements at the Heisenberg limit. Alternatively, if $\hat{H}$ is Hermitian, it is a quantum

observable, thus $\hat{H}^2$ would allow a direct estimate of variance, and so on. Unfortunately, the subspace $\mathcal{H}_G \cup \mathcal{H}_{G^\perp}$ is not invariant under $\hat{U}$ in general. As a result, repeated applications in this basis do not produce higher moments of $\hat{H}$, and instead leak out of $\mathcal{H}_G \cup \mathcal{H}_{G^\perp}$. The manner they do so is extremely difficult to predict, must be analyzed on a case-by-case basis, and thus is of limited utility.

Order can be restored to this undefined behavior by stemming the leakage. The simplest possibility that preserves the signal operator of Eq. 2 replaces $\hat{U}$ with a unitary, the *iterate*, that on an input in $\mathcal{H}_G \cup \mathcal{H}_{G^\perp}$ has the form

$$\hat{W} = \begin{pmatrix} \hat{H} & -g[\hat{H}] \\ g[\hat{H}] & \hat{H}^\dagger \end{pmatrix} = \sum_\lambda \begin{pmatrix} \lambda e^{i\theta_\lambda} & -\sqrt{1-\lambda^2} \\ \sqrt{1-\lambda^2} & \lambda e^{-i\theta_\lambda} \end{pmatrix} \otimes |\lambda\rangle\langle\lambda|, \quad \langle G|\hat{W}|G\rangle = \hat{H}. \tag{3}$$

Thus for every eigenstate of $\hat{H}$, $\hat{W}$ is isomorphic to SU(2). In the following, $\hat{W}$ will always be applied to states inside this subspace, thus its action on states outside it need not be defined. The usefulness of this construct is evident; due to its invariant subspace, multiple applications of the iterate result in highly structured behavior. However, implementing $\hat{W}$ requires $g[\hat{H}]$, which appears difficult to compute efficiently in general. In prior art [22], this was approximated using phase estimation. We formalize 'qubitization' as the problem of finding $\hat{W}$:

Qubitization: Given an m qubit signal oracle $\hat{U}$ and dimension d signal state $|G\rangle$ such that $\langle G|\hat{U}|G\rangle = \hat{H}$ is an n qubit normal signal operator, create a unitary $\hat{W} = \left(\begin{smallmatrix} \hat{H} & -g[\hat{H}] \\ g[\hat{H}] & \hat{H}^\dagger \end{smallmatrix}\right) \otimes \cdots$ where $\langle G|\hat{W}|G\rangle = \hat{H}$, and with an SU(2) invariant subspace containing $|G\rangle$.

Inputs: (1) controlled-$\hat{U}$ and its Hermitian conjugate; (2) operator $\hat{G}$ that prepares $|G\rangle$ from a standard basis state $\hat{G}|0\rangle = |G\rangle$ and its Hermitian conjugate; (3) $m+1$ qubits.
Output: $\hat{W} \in \mathbb{C}^{2^{m+1} \times 2^{m+1}}$.
Runtime: 1 query to inputs (1,2). $\mathcal{O}(1)$ primitive gates. Succeeds with probability 1.
Procedure: Thms. 2, 3 in Sec. 3.

Using the oracles $\hat{G}, \hat{U}$, and arbitrary unitary operations on *only* the ancilla register, we provide necessary and sufficient conditions, for the case of Hermitian $\hat{H}$ in Thm. 2 for when $\hat{W}$ can be implemented exactly using only one query to $\hat{U}$. We then prove in Thm. 3 that by using the *controlled-$\hat{U}$* oracle, there always exists a quantum circuit with the same signal operator and satisfies these conditions. We describe a similar construction for normal operators in Appendix. A.

When $\hat{H}$ is Hermitian, $\hat{W}^L$ produces Chebyshev polynomials $A[\hat{H}] = T_L[\hat{H}]$ [19]. We call any function $[\cdot]$ of the signal $\hat{H}$ *target operators* when they occur in the top-left block. The fact that Chebyshev polynomials are the best polynomial basis for L_∞ function approximation [37] suggests that the target operator $A[\hat{H}] + iB[\hat{H}]$ could approximate any desired function with a judicious choice of controls on the ancilla register. As all Hermitian operators are also quantum observables, achieving this would be of great utility to designing measurements on quantum states. We call this the *Observable transformations* problem.

Observable transformations: Given real functions $A, B : \mathbb{R} \to \mathbb{R}$, an m qubit signal oracle $\hat{U}$, and dimension d signal state $|G\rangle \in \mathbb{C}^d$ such that the n qubit signal operator $\langle G|\hat{U}|G\rangle = \hat{H}$ is Hermitian, create a unitary $\hat{W}_{\vec{\phi}}$ such that $\langle G|\hat{W}|G\rangle = A[\hat{H}] + iB[\hat{H}]$.

Inputs: (1) $\hat{W}$ obtained from qubitization of $\hat{U}$; (2) $\hat{G}, \hat{G}^\dagger$ from qubitization, (3) An integer $L > 0$. (2) Real polynomials $A(x), B(x)$ of degree L and parity L that satisfy the conditions of Thm. 5; (4) $m+1$ qubits.

Output: $\hat{W}_{\vec{\phi}} \in \mathbb{C}^{2^{m+1} \times 2^{m+1}}$.
Runtime: L queries to inputs (1) and $L+1$ queries to inputs (2). $\mathcal{O}(1)$ primitive gates. Succeeds with probability 1.
Procedure: Thm. 5 in Sec. 4.1.

The two polynomials in our solution of the Observable transformations problem are of the same parity. Polynomials of opposite parity can be obtained by embedding $\hat{W}$, or even $\hat{W}_{\vec{\phi}}$ into yet another SU(2) invariant subspace. We call this the *quantum signal processing* problem, which we addressed previously in [33].

Quantum signal processing: Given a real function $h : \mathbb{R} \to \mathbb{R}$ and an m qubit oracle $\hat{V} = \sum_\lambda e^{i\theta_\lambda} |\lambda\rangle\langle\lambda|$, create a unitary $\hat{V}_{\text{ideal}} = \sum_\lambda e^{ih(\theta_\lambda)} |\lambda\rangle\langle\lambda|$

Inputs: (1) Controlled-$\hat{V}$ and its Hermitian conjugate; (2) an odd periodic function $h : (-\pi, \pi] \to (-\pi, \pi]$; (3) target error $\epsilon > 0$ in trace distance; (4) $m+1$ qubits.
Output: $\hat{V}_{\vec{\phi}} \in \mathbb{C}^{2^{m+1} \times 2^{m+1}}$ such that $\|\langle +|\hat{V}_{\vec{\phi}}|+\rangle - \hat{V}_{\text{ideal}}\| \leq \epsilon$
Runtime: L total queries to inputs (1), where $L/2$, such that $\min_L \min_{A,C} |A[\theta] + iC[\theta] - e^{ih(\theta)}| \leq \epsilon/8$, is order of the best Fourier approximation to $e^{ih(\theta)}$, with A, C being Fourier series in θ of order $L/2$. $\mathcal{O}(1)$ primitive gates. Succeeds with probability $\geq 1 - 2\epsilon$.
Procedure: Thm. 6 in Sec. 4.2.

These schemes for target operator processing would be augmented by any constructive approach to implementing some $\hat{G}, \hat{U}$ consistent with a desired $\hat{H}$ – the linear-combination-of-unitaries algorithm [30], reviewed in Sec. 3.1, is one such possibility. In fact, as any complex operator $\hat{H} = \sum_j^d \alpha_j \hat{U}_j$ may be decomposed into a linear combination of unitaries, we conjecture that this combined with our algorithm for Observable transformations is an optimal approach to implementing arbitrary quantum measurements and their operator transformations. This claim is supported by its query complexity matching fundamental lower bounds in polynomial approximation theory, and our main result on Hamiltonian simulation.

We consider a general formulation of Hamiltonian simulation, outlined below. By applying our methods for target operator processing, an algorithm optimal in query complexity for all parameters values, and not just in asymptotic limits, is found.

Hamiltonian simulation: Given an m qubit signal oracle $\hat{U}$ and dimension d signal state $|G\rangle$ such that $\langle G|\hat{U}|G\rangle = \hat{H}$ is an n qubit Hamiltonian, create a unitary $\langle G|\hat{W}_{\text{ideal}}|G\rangle = e^{-i\hat{H}t}$ for some $t \in \mathbb{R}$.

Inputs: (1) Controlled-$\hat{U}$ and its Hermitian conjugate; (2) simulation time $t > 0$; (3) target error $\epsilon > 0$ in trace distance; (4) $m+2$ qubits.
Output: $\hat{V}_{\vec{\phi}} \in \mathbb{C}^{2^{m+2} \times 2^{m+2}}$ such that $\|\langle G|(\langle +|\hat{V}_{\vec{\phi}}|+\rangle - \hat{W}_{\text{ideal}})|G\rangle\| \leq \epsilon$.
Runtime: $\mathcal{O}(t + \log(1/\epsilon))$ queries to inputs (1). $\mathcal{O}(1)$ primitive gates. Succeeds with probability $\geq 1 - \mathcal{O}(\epsilon)$.
Procedure: Thm. 1 and Sec. 5. See Cor. 7, 8, 9 for applications.

Having motivated our perspective on the problem of qubitization in a quantum signal processor and its significance to designing arbitrary operator transformations, we present the solution.

3 Qubitization in a Quantum Signal Processor

This section describes, in three steps, qubitization: the process for creating the iterate $\hat{W}$ and an essential component in a systematic procedure for implementing operator transformations of $\hat{H}$. In Thm. 2, we provide necessary and sufficient conditions on when $\hat{W}$ can be constructed from the oracles $\hat{G}, \hat{U}$. Then in Thm. 3, we show that any $\hat{G}, \hat{U}$ not satisfying these conditions can be efficiently transformed into a $\hat{G}', \hat{U}'$ that do, and encode the same signal operator $\hat{H} = \langle G'|\hat{U}'|G'\rangle$. A constructive approach to implementing an encoding $\hat{U}$ for any desired $\hat{H}$ is then reviewed in Sec. 3.1.

For now, we assume that $|G\rangle$ is known. This soon proves to be unnecessary and only oracle access to $\hat{G}$ is required. Thus we must find a unitary $\hat{S}'$, acting only on the ancilla register such that the iterate $\hat{W} = \hat{S}'\hat{U}$ of Eq. 3 is obtained. For the case of Hermitian $\hat{H}|\lambda\rangle = \lambda|\lambda\rangle$, we now determine necessary and sufficient conditions on what $\hat{S}'$ must be. As $\hat{S}'$ is otherwise arbitrary, we use without loss of generality the ansatz of $\hat{S}'$ being a product of a reflection about $|G\rangle$ and another arbitrary unitary $\hat{S}$ on the ancilla:

$$\hat{W} = (2|G\rangle\langle G| - \hat{I})\hat{S}\hat{U}, \quad |G_\lambda\rangle = |G\rangle|\lambda\rangle \Rightarrow |G_\lambda^\perp\rangle = \frac{\lambda|G_\lambda\rangle - \hat{S}\hat{U}|G_\lambda\rangle}{\sqrt{1-\lambda^2}}. \tag{4}$$

Theorem 2 (Conditions on Hermitian qubitization)**.** *For all signal oracles $\hat{U}$ that implement the signal operator $\hat{H}$, the unitary $\hat{S}$ in Eq. 4 creates a unitary iterate $\hat{W}$ with the same signal operator in the same basis, but in an $SU(2)$ invariant subspace containing $|G\rangle$ if and only if*

$$\langle G|\hat{S}\hat{U}|G\rangle = \hat{H} \text{ and } \langle G|\hat{S}\hat{U}\hat{S}\hat{U}|G\rangle = \hat{I}. \tag{5}$$

Proof. In the forward direction, we assume Eq. 3, then compute and compare $\langle G_\lambda|\hat{W}|G_\lambda\rangle = \langle G_\lambda|\hat{S}\hat{U}|G_\lambda\rangle = \lambda$. By Gram-Schmidt orthonormalization of $\hat{W}|G_\lambda\rangle$, we also obtain the state $|G_\lambda^\perp\rangle = \frac{\lambda|G_\lambda\rangle - \hat{S}\hat{U}|G_\lambda\rangle}{\sqrt{1-\lambda^2}}$ orthogonal to $|G_\lambda\rangle$. Thus we compute and compare $\langle G_\lambda|\hat{W}|G_\lambda^\perp\rangle = \frac{\lambda^2 - \langle G_\lambda|(\hat{S}\hat{U})^2|G_\lambda\rangle}{\sqrt{1-\lambda^2}} = -\sqrt{1-\lambda^2} \Rightarrow \langle G_\lambda|(\hat{U}\hat{S})^2|G_\lambda\rangle = 1$. As these must be true for all eigenvectors $|\lambda\rangle$, the conditions in Eq. 5 are necessary. That these are also sufficient follows from assuming Eq. 5 and a straightforward computation of $\hat{W}|G_\lambda\rangle$ and $\hat{W}|G_\lambda^\perp\rangle$, which recovers Eq. 3. □

In hindsight, these results are manifest. After all, $\langle G|\hat{S}\hat{U}\hat{S}\hat{U}|G\rangle = \hat{I}$ implies that $\hat{S}\hat{U}$ is a reflection when controlled by input state $|G\rangle$, and it is well-known that a Grover iterate [25, 45] is the product of two reflection about start and target subspaces. Nevertheless, the sufficiency of these conditions highlights that this is the simplest method to extract controllable and predictable behavior out of $\hat{U}$.

Unfortunately, depending on the details of $\hat{U}$, a solution to Eq. 5 may not exist. Thm. 2, amounts to choosing $\hat{S}$ such that $\hat{S}\hat{U}\hat{S}$ is the inverse $\hat{U}^\dagger$ whist preserving the signal operator $\langle G|\hat{S}\hat{U}|G\rangle = \hat{H}$. Given that $\hat{S}$ only acts on the ancilla register, it is hard to see how this is possible in general. The solution is to construct a different quantum circuit $\hat{U}'$ that contains $\hat{U}$ but still implements the same signal operator, and crucially always has a solution $\hat{S}$. We now show how this can be done in all cases using only 1 query to controlled-$\hat{U}$ and controlled-$\hat{U}^\dagger$.

Theorem 3 (Existence of Hermitian qubitization)**.** *For all m qubit signal unitaries $\hat{U}$ that implement the signal operator $\langle G|\hat{U}|G\rangle = \hat{H}$, there exists an $m+1$ qubut quantum circuit $\hat{U}'$ that queries controlled-$\hat{U}$ and controlled-$\hat{U}^\dagger$ once to implement the Hermitian component $\frac{1}{2}(\hat{H} + \hat{H}^\dagger)$ as the signal operator, such that the of conditions Eq. 5 can be satisfied.*

Proof. We prove this by an explicit construction. Let the controlled-$\hat{U}$ operators be $\hat{Q}_1 = |0\rangle\langle 0| \otimes \hat{I} + |1\rangle\langle 1| \otimes \hat{U}^\dagger$, $\hat{Q}_2 = |0\rangle\langle 0| \otimes \hat{U} + |1\rangle\langle 1| \otimes \hat{I}$. Thus the extra qubit states $|0\rangle, |1\rangle$, are flags that selects either $\hat{U}_j$ or $\hat{U}_j^\dagger$. By multiplying, $\hat{U}' = \hat{Q}_1\hat{Q}_2 = |0\rangle\langle 0| \otimes \hat{U} + |1\rangle\langle 1| \otimes \hat{U}^\dagger$. Now consider the ancilla state $|G'\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|G\rangle$, and choose $\hat{S} = (|0\rangle\langle 1| + |1\rangle\langle 0|) \otimes \hat{I}_{as}$. It is easy to verify that the conditions of Eq. 5 is satisfied.

$$\langle G'|\hat{S}\hat{U}'|G'\rangle = \langle G'|\hat{U}'|G'\rangle = \frac{1}{2}(\hat{H} + \hat{H}^\dagger), \quad \langle G'|\hat{S}\hat{U}'\hat{S}\hat{U}'|G'\rangle = \langle G'|\hat{U}'^\dagger\hat{U}'|G'\rangle = \hat{I} \tag{6}$$

where we have used the fact that $\hat{S}|G'\rangle = |G'\rangle$ is an eigenstate, and that $\hat{S}$ swaps the $|0\rangle, |1\rangle$ ancilla states in $\hat{U}'$, thus transforming it into its inverse. □

Even if we are given $\hat{U}$ for which there is no solution to Eq. 5, we can always apply Thm. 3 to construct a $\hat{U}'$ that does with minimal overhead. Furthermore our proof uses no information about the detailed structure of $|G\rangle$. Thus without loss of generality, we can assume that any $\hat{G}, \hat{U}$ have already been qubitized.

3.1 Implementing $\hat{G}, \hat{U}$ with a Linear Combination of Unitaries

Qubitization can be made fully constructive with an approach for implementing some $\hat{G}, \hat{U}$ that encode any desired $\hat{H}$. One option is provided by the Linear-Combination-of-Unitaries algorithm (LCU) [20, 30], which underlies the BCCKS simulation algorithm. LCU is based on the fact that any complex $\hat{H}$ is a linear combination of some d unitary operators:

$$\hat{H} = \sum_{j=1}^{d} \alpha_j \hat{U}_j, \quad \|\hat{H}\| \leq \sum_{j=1}^{d} |\alpha_j| = \alpha, \tag{7}$$

where the upper bound on the spectral norm is α. Note that this bound depends on the choice of decomposition, but is tight for the some choice. Without loss of generality, all $\alpha_j \geq 0$ by absorbing complex phases into $\hat{U}_j$. The algorithm assumes that the α_j are provided as a list of d numbers, and each $\hat{U}_j$ is provided as a quantum circuit composed of $\mathcal{O}(C)$ primitive gates. With these inputs, the oracles

$$\hat{G} = \sum_{j=1}^{d} \sqrt{\frac{\alpha_j}{\alpha}} |j\rangle\langle 0|_a + \cdots, \quad \hat{U} = \sum_{j=1}^{d} |j\rangle\langle j|_a \otimes \hat{U}_j, \quad \langle G|\hat{U}|G\rangle = \frac{\hat{H}}{\alpha} \tag{8}$$

can be constructed, where the ancilla state creation operator $\hat{G}|0\rangle = |G\rangle$ is implemented with $\mathcal{O}(d)$ primitive gates, and the selector $\hat{U}$ is implemented with $\mathcal{O}(dC)$ primitive gate. By direct expansion of $\hat{U}\hat{G}|0\rangle_a|\psi\rangle_s$, this leads exactly to Eq. 2. Of course, the optimal decomposition that costs the fewest number of ancilla qubits and primitive gates may be difficult to find, and may not even fit naturally in this model, but LCU shows that implementing an encoding for any $\hat{H}$ is possible in principle.

4 Operator Function Design on a Quantum Signal Processor

The purpose of the quantum signal processor is to transform the signal $\hat{H}$ into any desired target operator $f[\hat{H}]$ – the observable transformations and quantum signal processing problems. We present a systematic framework that furnishes the optimal complexity and a concrete procedure for

almost any f and show how an exact connection is made between query complexity and the theory of best function approximations with polynomials [36, 37].

Qubitization in Sec. 3 is the essential first step that makes this endeavor plausible, as evidenced by the highly structured behavior of the iterate $\hat{W}$ in Eq. 3, where for Hermitian $\hat{H}$, multiple applications elegantly generate Chebyshev polynomials $T_L[\hat{H}]$ [19]. To go further, additional control parameters on $\hat{W}$ are necessary, and in the following, we only consider Hermitian $\hat{H}$. Thus we introduce the *phased iterate* with the same invariant subspace as $\hat{W}$:

$$\hat{W}_\phi = \begin{pmatrix} \hat{H} & -ie^{-i\phi}g[\hat{H}] \\ -ie^{i\phi}g[\hat{H}] & \hat{H}^\dagger \end{pmatrix}, \quad \langle G|\hat{W}_\phi|G\rangle = \hat{H}. \tag{9}$$

Lemma 4. *The phased iterate $\hat{W}_\phi = \hat{Z}_{\phi-\pi/2}\hat{W}\hat{Z}_{-\phi+\pi/2}$, where $\hat{Z}_\phi = ((1+e^{-i\phi})|G\rangle\langle G| - \hat{I})$ is a partial reflection about $|G\rangle$ by angle $\phi \in \mathbb{R}$, implements a relative phase between the $|G_\lambda\rangle$ and $|G_\lambda^\perp\rangle$ subspaces. In block form,*

$$\hat{Z}_\phi = e^{-i\phi/2}\sum_\lambda \begin{pmatrix} e^{-i\phi/2} & 0 \\ 0 & e^{i\phi/2} \end{pmatrix} \otimes |\lambda\rangle\langle\lambda|. \tag{10}$$

Proof. By direct computation. □

We provide algorithms for two large classes of transformations. 'Observable transformations' in Sec. 4.1 implement target operators where the real and imaginary parts have the same parity with respect to $\hat{H}$. This is complemented by 'quantum signal processing' in Sec. 4.2 for target operators with opposite parity.

4.1 Observable Transformations

A sequence of L of phased iterates is a product of SU(2) rotations

$$\hat{W}_{\vec{\phi}} = \hat{W}_{\phi_L}...\hat{W}_{\phi_2}\hat{W}_{\phi_1} = \hat{I}_2 \otimes A[\hat{H}] + i(\hat{\sigma}_z \otimes B[\hat{H}] + \hat{\sigma}_x \otimes C[\hat{H}] + \hat{\sigma}_y \otimes D[\hat{H}]), \tag{11}$$

where $\hat{I}_2, \hat{\sigma}_{x,y,z}$ are the identity and Pauli matrices respectively in the $|G_\lambda\rangle$ and $|G_\lambda^\perp\rangle$ basis, and (A, B, C, D) are operator functions of $\hat{H}$. As we can only prepare and measure states spanned by the $|G_\lambda\rangle$, the observable transformations problem is concerned with the component $\langle G|\hat{W}_{\vec{\phi}}|G\rangle = A[\hat{H}] + iB[\hat{H}]$. Any choice of phases $\vec{\phi} \in \mathbb{R}^L$ generates sophisticated interference effects between elements of the sequence, leading to (A, B, C, D) with some non-trivial functional dependence on $\hat{H}$. Though the dependence of the output on $\vec{\phi}$ seems hard to intuit, they nevertheless specify a *program* for computing functions of $\hat{H}$, similar to how a list of numbers might specify a polynomial.

Remarkably, one can specify the desired output $A[\hat{H}] + iB[\hat{H}]$, and then efficiently invert this specification to obtain its list of $\vec{\phi}$. Clearly not all choices are achievable, but the space of possibilities turns out to be large.

Theorem 5 (Achievable Observable Transformations)**.** *A choice of functions (A, B) in Eq. 11 can be implemented by some $\vec{\phi} \in \mathbb{R}^L$ if and only if all the following are true: (1) $A(x), B(x)$ are real degree L polynomials of parity L; (2) $A(x) = 1$; (3) $\forall|x| \leq 1,\ A^2(x) + B^2(x) \leq 1$; (4) $\forall|x| \geq 1,\ A^2(x) + B^2(x) \geq 1$; (5) $\forall L$ even, $|x| \geq 1,\ A^2(ix) + B^2(ix) \geq 1$.*

Proof. Previously [35], we studied the product of single qubit rotations $e^{-i\theta\hat{\sigma}_{\phi_L}/2}...e^{-i\theta\hat{\sigma}_{\phi_2}/2}e^{-i\theta\hat{\sigma}_{\phi_1}/2} = \hat{I}_2 A + i\left(\hat{\sigma}_z B + \hat{\sigma}_x C + \hat{\sigma}_y D\right)$, where $\hat{\sigma}_\phi = \cos(\phi)\hat{\sigma}_x + \sin(\phi)\hat{\sigma}_y$, provided a full characterization of possible (A, B, C, D), and an algorithm to efficiently compute $\vec{\phi}$ from any partial specification of at most two functions. These are isomorphic to the SU(2) block of Eq. 11, thus the results carry over. Specifically, Thm. 5 follows from Thms. 1,2 of [35]. □

As polynomials form a complete basis on bounded real intervals, these results imply the query complexity of approximating any real function $A[\hat{H}]$ with error ϵ is exactly that of its best polynomial ϵ-approximation satisfying the constraints of Thm. 5, and similarly for the complex case.

4.2 Quantum Signal Processing

The solution to observable transformations in Thm. 5 implements a function of $\hat{H}$ where the real and complex parts have the same parity. We now describe how complementary behavior where these parts of the target operator have opposite parity. The main results of this section are drawn from [33] and are included for completeness. Observe that the SU(2) invariant subspace of the iterate times a global phase $e^{i\Phi}\hat{W}_\phi$ can be diagonalized to obtain

$$e^{i\Phi}\hat{W}_\phi = e^{i\Phi - i\hat{\sigma}_\phi \otimes \cos^{-1}[\hat{H}]}, \quad \hat{\sigma}_\phi = \cos(\phi)\hat{\sigma}_x + \sin(\phi)\hat{\sigma}_y \tag{12}$$

with eigenvectors $|G_{\lambda\pm}\rangle = \frac{|G_\lambda\rangle \pm e^{i\phi}|G_\lambda^\perp\rangle}{\sqrt{2}}$. Quantum signal processing allows us to approximate

$$\hat{W}_{\text{ideal}} = e^{ih[\Phi - \hat{\sigma}_\phi \otimes \cos^{-1}[\hat{H}]]}, \tag{13}$$

where h is any odd real periodic function. This is implemented by a sequence $\langle +|_b \hat{V}_{\vec{\varphi}} |+\rangle_b$ of controlled-$\hat{W}$ operators $\hat{V}_0 = |+\rangle\langle +|_b \otimes \hat{I}_{as} + |-\rangle\langle -|_b \otimes (e^{i\Phi}\hat{W}_\phi)$, where $\hat{\sigma}_x|\pm\rangle = \pm|\pm\rangle$, that acts on and is then projected on the state $|+\rangle_b$ of the new ancilla qubit:

$$\hat{V}_{\vec{\varphi}} = \hat{V}_{\varphi_N}^\dagger ... \hat{V}_{\varphi_4+\pi}^\dagger \hat{V}_{\varphi_3} \hat{V}_{\varphi_2+\pi}^\dagger \hat{V}_{\varphi_1}, \quad \hat{V}_\varphi = (e^{-i\varphi\hat{\sigma}_z/2} \otimes \hat{I}_{as})\hat{V}_0(e^{i\varphi\hat{\sigma}_z/2} \otimes \hat{I}_{as}). \tag{14}$$

Similar to Thm. 5, a large class of eigenphase functions can be approximated, but with one complication. Up to this point, all our constructions have been deterministic. In contrast, Quantum signal processing requires a probabilistic projection onto an ancilla state. However, this is not problematic as the failure probability can be made arbitrarily small.

The query complexity of implementing this approximation is given by

Theorem 6 (Quantum Signal Processing [33])**.** *$\forall$ real odd periodic functions $h : (-\pi, \pi] \to (-\pi, \pi]$ and even $N > 0$, let $(A[\theta], C[\theta])$ be real Fourier series in $(\cos(k\theta), \sin(k\theta))$, $k = 0, ..., N/2$, that approximates $\max_{\theta\in\mathbb{R}} |A[\theta] + iC[\theta] - e^{ih(\theta)}| \leq \epsilon/8$. Given $A[\theta], C[\theta]$, one can efficiently compute the $\vec{\phi}$ such that $\langle +|_b \hat{V}_{\vec{\varphi}} |+\rangle_b$ in Eq. 14 applies $\hat{V}_\phi$ a number N times to approximate $\hat{W}_{ideal}$ in Eq. 13 with trace distance $\|\langle +|_b \hat{V}_{\vec{\varphi}} |+\rangle_b - \hat{W}_{ideal}\| \leq \epsilon$ and success probability $\geq 1 - 2\epsilon$.*

This result complements Thm. 5 as the Fourier basis is also complete for periodic functions.

5 Application to Hamiltonian Simulation

With these results for qubitization and operator function design with a quantum signal processor, the application to Hamiltonian simulation follows easily. We complete the proof Thm. 1, and then apply these results to obtain our claims of improvements in Table. 1.

Suppose we are given access to a unitary $\hat{G}$ that prepares state $|G\rangle = \hat{G}|0\rangle$ from some standard computational basis state $|0\rangle$, and a unitary $\hat{U}$ on q qubits such that $\langle G|\hat{U}|G\rangle = \hat{C}$ implements a complex target operator. By qubitizing $\hat{G}, \hat{U}$ with an additional qubit in Thm. 3, we obtain a modified $\hat{U}'$ with signal operator $\hat{H} = \frac{1}{2}(\hat{C} + \hat{C}^\dagger)$. This $\hat{U}'$ has an iterate $\hat{W}$ in Eq. 3 which can be multiplied by a global phase $e^{i\Phi}$ to obtain $e^{i\Phi}\hat{W} = e^{i\Phi - i\hat{\sigma}_y \otimes \cos^{-1}(\hat{H})}$ in Eq. 12 to implement a

nonlinear function of $\hat{H}$ on its eigenstates. Hamiltonian simulation is accomplished by linearization this phase with a suitable choice of Φ and h with yet another qubit in Thm. 6 such that

$$h[\Phi - \hat{\sigma}_y \otimes \cos^{-1}[\hat{H}]] = -\hat{I}_2 \otimes \hat{H}t \quad \Rightarrow \quad \langle G|\hat{W}_{\text{ideal}}|G\rangle = e^{-i\hat{H}t}. \tag{15}$$

As done in [10, 33], it is easily verified that one suitable choice is

$$\Phi = -\pi/2, \quad h(\theta) = \sin(\theta)t. \tag{16}$$

Thm. 6 requires an $N/2$ order Fourier approximation to $e^{i\sin(\theta)t}$, which can be obtained by the truncating its Jacobi-Anger expansion $e^{i\sin(\theta)t} = \sum_{k=-\infty}^{\infty} J_k(t)e^{ikt}$ [2], with error [10]

$$\epsilon \leq \sum_{k=q}^{\infty} 2|J_k(t)| \leq \frac{4t^q}{2^q q!} = \mathcal{O}\Big(\Big(\frac{et}{2q}\Big)^q\Big), \quad q = 1 + \frac{N}{2} \quad \Rightarrow \quad \log\left(\frac{1}{\epsilon}\right) = \mathcal{O}\left(q\log\left(\frac{2q}{e\tau}\right)\right), \tag{17}$$

where $J_k(t)$ is the k^{th} Bessel function of the first kind. The rapid converge by truncation arises as $e^{it\sin(\theta)}$ is an entire analytic function [14]. Solving for N [33] then furnishes the number of queries to $\hat{W}$ required to simulate $e^{-i\hat{H}t}$ with error $\mathcal{O}(\epsilon)$ and failure probability $\mathcal{O}(\epsilon)$:

$$N = \mathcal{O}(t + \gamma(t,\epsilon)), \quad \gamma(t,\epsilon) = \begin{cases} \mathcal{O}(\log(1/\epsilon)/\log\log(1/\epsilon)), & t = \mathcal{O}(\gamma(t,\epsilon)), \\ \mathcal{O}(\log(1/\epsilon)), & \text{otherwise.} \end{cases} \tag{18}$$

This achieves the upper bound in Thm. 1. The tradeoff between ϵ, t, N in Eq. 17 is plotted in Appendix B, together with example phases $\vec{\varphi}$ implementing Eq. 16. To prove that it is optimal, we show that sparse Hamiltonian simulation is a special case.

Corollary 7 (Hamiltonian Simulation of a Sparse Hermitian Matrix)**.** *Given access to oracle $\hat{O}_H|j\rangle|k\rangle|z\rangle = |j\rangle|k\rangle|z \oplus \hat{H}_{j,k}\rangle$ queried by $j \in [2^n]$ row and $k \in [2^n]$ column indices to return the value $\hat{H}_{j,k} = \langle j|\hat{H}|k\rangle$, with maximum absolute value $\|\hat{H}\|_{max} = \max_{jk}|\hat{H}_{j,k}|$, and oracle $\hat{O}_F|j\rangle|l\rangle = |j\rangle|f(j,l)\rangle$ queried by $j \in [2^n]$ row and $l \in [d]$ column indices to compute in-place the column index $f(j,l)$ of the l^{th} non-zero entry of the j^{th} row, time evolution by $\hat{H}$ can be simulated with for time t and error ϵ with $\mathcal{O}(dt\|\hat{H}\|_{max} + \gamma(t,\epsilon))$ queries to $\hat{O}_H, \hat{O}_F$.*

Proof. Let $|G\rangle = |0\rangle_{a_1}|0\rangle_{a_2}$. Using 2 queries to $\hat{O}_H$ and 1 query to $\hat{O}_F$, [11] shows how the isometry $\hat{T}_1 = \sum_{j\in[2^n]}|\psi_j\rangle\langle 0|_{a_1}\langle 0|_{a_2}\langle j|_s$, $\hat{T}_2 = \sum_{j\in[2^n]}|\chi_j\rangle\langle 0|_{a_1}\langle 0|_{a_2}\langle j|_s$ each can be implemented, where $|\psi_j\rangle = \sum_{l\in[d]}\frac{|f(j,l)\rangle_{a_2}}{\sqrt{d}}\Big(\sqrt{\frac{H_{j,f(j,l)}}{\|\hat{H}\|_{\max}}}|0\rangle_{a_1} + \sqrt{1 - \frac{|H_{j,f(j,l)}|}{\|\hat{H}\|_{\max}}}|1\rangle_{a_1}\Big)|j\rangle_s$ and $|\chi_j\rangle = \sum_{l\in[d]}\frac{|f(j,l)\rangle_s}{\sqrt{d}}\Big(\sqrt{\frac{H^*_{j,f(j,l)}}{\|\hat{H}\|_{\max}}}|0\rangle_{a_1} + \sqrt{1 - \frac{|H_{j,f(j,l)}|}{\|\hat{H}\|_{\max}}}|1\rangle_{a_1}\Big)|j\rangle_{a_2}$. Let $\hat{U} = \hat{T}_2^\dagger\hat{T}_1$. The lemma is proven from the upper bound of Thm. 1 if $\langle G|\hat{U}|G\rangle = \hat{H}/d\|\hat{H}\|_{\max}$. This is true by direct computation. □

The query complexity in Eq. 17 for this sparse case exactly matches a lower bound based on simulating a Hamiltonian that solves PARITY with unbound error [8], valid for all parameter values, and not just in asymptotic limits [10, 33]. This completes the proof of Thm. 1.

The case where $\hat{H}$ decomposes into a linear combination of unitaries is an immediate application:

Corollary 8 (Hamiltonian Simulation of a Linear Combination of Unitaries)**.** *Given access to an oracle $\hat{G}$ that prepares $|G\rangle = \sum_{j=1}^{d}\sqrt{\alpha_j/\alpha}|j\rangle_a$, where $\alpha_j \geq 0$, $\alpha = \sum_{j=1}^{d}\alpha_j$ and oracle $\hat{U} = \sum_{j=i}^{d}|j\rangle\langle j|_a \otimes \hat{U}_j$, time evolution by $\hat{H} = \sum_{j=1}^{d}\alpha_j\hat{U}_j$ can be simulated with for time t and error ϵ with $\mathcal{O}(\alpha t + \gamma(t,\epsilon))$ queries to $\hat{G}, \hat{U}$.*

Proof. The lemma is proven from Thm. 1 if $\langle G|\hat{U}|G\rangle = \hat{H}/\alpha$. This is true by direct computation. □

An example of this algorithm this is considered in Appendix. C for simulating the dissociation of molecular hydrogen at chemical accuracy.

The intuitiveness of Thm. 1 allows us to swiftly devise new models of Hamiltonian simulation.

Corollary 9 (Hamiltonian Simulation of a Purified Density Matrix)**.** *Given access to an oracle $\hat{G}$ that prepares a purification $\hat{G}|0\rangle_a = |G\rangle_a = \sum_j \sqrt{\alpha_j}|j\rangle_{a_1}|\chi_j\rangle_{a_2}$ of density matrix $\hat{\rho} = Tr_{a_1}[|G\rangle\langle G|]$, time evolution by $\hat{\rho}$ can be simulated with for time t and error ϵ with $\mathcal{O}(t + \gamma(t, \epsilon))$ queries to $\hat{G}$.*

Proof. The lemma is proved if we can find $\hat{U}$ such that Thm. 1 $\langle G|\hat{U}|G\rangle = \hat{\rho}$. We choose a $\hat{U}$ that swaps the system s and ancilla register a_2. Let $\{|\lambda\rangle\}$ be a complete basis on the system. By direct computation,

$$\langle G|\hat{U}|G\rangle_a \sum_\lambda |\lambda\rangle\langle\lambda|_s = \sum_\lambda \sum_j \langle G|\sqrt{\alpha_j}|j\rangle_{a_1}|\lambda\rangle_{a_2}|\chi_j\rangle\langle\lambda|_s = \sum_\lambda \sum_j |\alpha_j||\chi_j\rangle_s\langle\chi_j|\lambda\rangle\langle\lambda|_s = \hat{\rho}. \quad (19)$$

□

6 Conclusion

Our general procedure for Hamiltonian simulation in Thm. 1 extends the scope of possible useful formulations of Hamiltonian simulation. As seen in Table. 1, it encompasses any case where the Hamiltonian is embedded in a flagged subspace of the signal unitary. Given this, a simulation algorithm with query complexity optimal in all parameters, and also not just in asymptotic limits, is easily obtained with minimal overhead. While this procedure contains and significantly improves upon important models where the Hamiltonian is d-sparse or a linear combination of unitaries, its greater value lies in illuminating an intuitive and straightforward path to other as-yet undiscovered models of Hamiltonian simulation. In particular, our result for time-evolution by a purified density matrix is a quadratic improvement in time and an exponential improvement in error over the sample-based model – the proof of which consisted of just a few lines.

Many other exciting directions extend from this work. One example is how additional structural information about $\hat{H}$ [18] may be exploited. This is illustrated by when the spectral norm of $\|\hat{H}\|$ is smaller than the sum of coefficients α of a particular linear combination of unitaries decomposition. If this decomposition were to be used, simulation would take time $\mathcal{O}(\alpha\|\hat{H}\|t+\log(1/\epsilon))$ – a factor α slowdown. In principle, an $\alpha = \mathcal{O}(1)$ decomposition always exists, but this may be difficult to find. Furthermore, it is easy to construct pathological $\hat{H}$ with small norms, but nevertheless decompose by naive methods into components with large spectral norms [42]. Our approach offers a possible solution – one finds any projector for the Hamiltonian, rather than some specific decomposition into parts with properties dictated by the formulation. It remains an interesting challenge to identify the cases where, and determine how structural information may be incorporated.

These advances are special cases arising from our vision of the more general quantum signal processor. Through qubitization, structure is imposed onto any unitary process implementing some Hermitian signal operator. This structure allows for efficient processing of the signal, by the techniques of observable transformations and quantum signal processing, into some more desired form. As the query complexity of approximating any arbitrary target function of the signal exactly matches fundamental bounds lower in polynomial and Fourier approximation theory [37], we expect this to have numerous application in metrology [34] and other quantum algorithms [45].

7 Acknowledgements

G.H.Low and I.L.Chuang thank Robin Kothari for insightful discussions and acknowledge funding by the ARO Quantum Algorithms Program and NSF RQCC Project No.1111337.

A Qubitization of Normal Operators

The results of Thms. 3, 2 for qubitization can be extended to normal operators. It is well known that any normal matrix has a polar decomposition

$$\hat{H} = \hat{H}_U \hat{H}_H, \tag{20}$$

where $\hat{H}_U$ is unitary, $\hat{H}_H$ is positive-semidefinite, and $[\hat{H}_U, \hat{H}_H] = 0$ commute, with eigenvalues of $\hat{H}|\lambda\rangle = e^{i\theta_\lambda}\lambda|\lambda\rangle$, where $\lambda \geq 0$, $\theta_\lambda \in \mathbb{R}$. This reduces to a Hermitian operator when $\hat{H}_U$ has eigenvalues ± 1, and reduces to a unitary operator when all $\hat{H}_H$ has eigenvalues 1. The trivial approach to qubitization applies to any complex matrix. We simply use the construction of Thm. 3 to implement the Hermitian target operator $\frac{1}{2}(\hat{H} + \hat{H}^\dagger)$.

Another possibility uses two phased iterates in an alternating sequence on input state $|G\rangle_a|\psi\rangle_s$,

$$\hat{W}_{\phi\pm} = \hat{Z}_{-\phi-\pi/2}(2|G\rangle\langle G| - \hat{I})\hat{U}_\pm \hat{Z}_{\phi+\pi/2}, \quad \hat{W}_{\vec{\phi}\pm} = \hat{W}_{\vec{\phi}} = \hat{W}_{\phi_L+}...\hat{W}_{\phi_3+}\hat{W}_{\phi_2-}\hat{W}_{\phi_1+}, \tag{21}$$

where $\hat{U}_+ = \hat{U}$ and $\hat{U}_- = \hat{U}^\dagger$. For each eigenstate $|\lambda\rangle$, the separate iterates have the block form

$$\hat{W}_{\phi+} = \begin{pmatrix} \lambda e^{i\theta_\lambda} & \cdot & -ie^{i\phi}\sqrt{1-\lambda^2} & \cdot \\ -ie^{-i\phi}\sqrt{1-\lambda^2} & \cdot & \lambda e^{-i\theta_\lambda} & \cdot \\ 0 & \cdot & 0 & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{pmatrix}, \quad \hat{W}_{\phi-} = \begin{pmatrix} \lambda e^{-i\theta_\lambda} & -ie^{-i\phi}\sqrt{1-\lambda^2} & \cdot & \cdot \\ 0 & 0 & \cdot & \cdot \\ -ie^{i\phi}\sqrt{1-\lambda^2} & \lambda e^{i\theta_\lambda} & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{pmatrix}, \tag{22}$$

where the first column corresponds to input states $\{|G\rangle|\lambda\rangle, |G\rangle\hat{V}_-|\lambda\rangle, |G\rangle\hat{V}_+|\lambda\rangle\}$. The subspace spanned by these states is not invariant under any repeated application of an iterate of the same sign. However, the product $\hat{W}_{\phi_i-}\hat{W}_{\phi_j+}$ has an invariant subspace containing $|G\rangle$. With the understanding that we only consider alternating sequences, each $\hat{W}_{\phi\pm}$ has the representation

$$\hat{W}_{\phi\pm} = \begin{pmatrix} e^{\pm i\theta_\lambda}\lambda & -ie^{-i\phi}\sqrt{1-\lambda^2} \\ -ie^{i\phi}\sqrt{1-\lambda^2} & e^{\mp i\theta_\lambda}\lambda \end{pmatrix}. \tag{23}$$

Note that when all eigenvalues λ are identical, this reduces to Oblivious amplitude amplification [7], and we recover Hermitian qubitization when all $\theta_\lambda = 0$. While this approach uses one less ancilla qubit than the construction of Thm. 3, quantum signal processing can only be performed on controlled block of even length $\hat{W}_{\vec{\phi}\pm}$, as only they have an invariant subspace. This limitation can be relevant in some cases, such as Hamiltonian simulation where quantum signal processing is applied to a single $\hat{W}_\phi$.

B Practical Details for Implementing Hamiltonian Simulation

This appendix illustrates a specific application of the quantum signal processing approach to a signal unitary that encodes the Hamiltonian $\hat{H}$ as a signal operator. In particular, a comparison of

performance with the BCCKS approach is made. The details will be useful to readers interested in implementing our procedure on a quantum computer. These include plots for the exact error scaling Eq. 17 of the $N/2$ term Fourier approximation to $e^{it\sin(\theta)}$ in Fig. 1(Left), and the number of required queries per unit of simulated time in Fig. 1(Right). Furthermore, a table of select phases $\vec{\varphi}$ computed using the algorithm in [35] can be found in Table. 3.

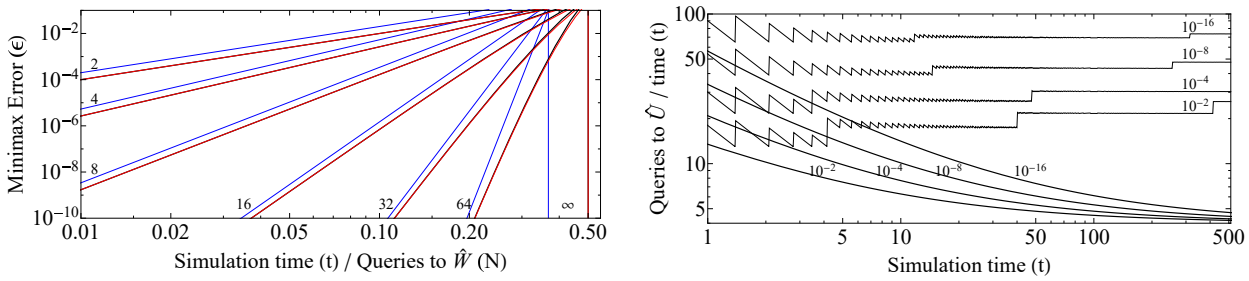

Figure 1: (Left) Approximation error $\epsilon = \max_{\theta\in\mathbb{R}} |A[\theta] - iC[\theta] - e^{it\sin(\theta)}|$. $(A[\theta], C[\theta])$ are real Fourier series in $(\cos(k\theta), \sin(k\theta))$, $k = 0, ..., N/2$, and ϵ is plotted for the upper bound $\frac{4t^q}{2^q q!}$ (blue), truncation $\sum_{k=q}^{\infty} 2|J_k[t]|$ (black), and best possible [27] (red), for $N = 2, 4, 8, 16, 32, 64, \infty$ queries to the controlled iterate $\hat{W}$, where $q = 1 + N/2$. (Right) Queries per unit of simulation time to unitary $\hat{U}$ encoding $\hat{H} = \langle G|\hat{U}|G\rangle$ at target trace distances $\|\langle +|\hat{V}_{\vec{\varphi}}|+\rangle - e^{-i\hat{H}t}\| = 10^{-2,-4,-8,-16}$ for the BCCKS algorithm (thin), and this work using the truncated approximation to $e^{it\sin(\theta)}$ (thick).

In the interests of a fair comparison, Fig. 1(Right) counts the number of queries to the signal operator $\hat{U}$ in the $\hat{H} = \langle G|\hat{U}|G\rangle$ encoding. $\hat{U}$ is not assumed to be qubitized, thus incurring a factor 2 additional cost from quering $\hat{U}$ and $\hat{U}^\dagger$ each over the asymptotic limit of 2 queries per unit of simulation time in Fig. 1(Left). Similarly, the BCCKS algorithm [9] incurs a factor 3 additional cost from querying $\hat{U}$ twice and $\hat{U}^\dagger$ in their use of oblivious amplitude amplification. As BCCKS is known to be optimal in the regime $t = \mathcal{O}(\frac{\log(t/\epsilon)}{\log\log(t/\epsilon)})$, the improvement of our approach is most dramatic outside of it. In particular, the queries per unit time of BCCKS scales like $\mathcal{O}(\frac{\log(t/\epsilon)}{\log\log(t/\epsilon)})$, whereas our approach approaches 4 in the limit $t \to \infty$.

The number of queries to in the BCCKS model is $3Kr$, where K is the queries per segment $e^{-i\hat{H}\log 2/r}$, $r = \lceil t/\log(2)\rceil$ is the number of segments, and 3 is for oblivious amplitude amplification on each segment. K is chosen such that $\sum_{k=K+1}^{\infty} \frac{\log^k 2}{k!} \leq \frac{\epsilon}{r}$. An analysis of the procedure [9] shows that the trace distance of its simulated evolution is a factor $\mathcal{O}(1) \gtrsim 2$ larger than ϵ. Thus Fig. 1 slightly overestimates BCCKS performance as we take ϵ directly to be the trace distance.

C Scalable Quantum Simulation of Molecular Energies

Recently, O'Malley et. al realized a breakthrough experimental quantum simulation of the dissociation of molecular hydrogen [39]. This featured the use of techniques believed essential for scaling quantum chemistry simulations to larger problem sizes. There, interactions between a minimal set of orbitals were mapped, through efficient classical pre-computation, to an effective Hamiltonian with $d+1$ terms $\hat{H}_j$ that captured the essential quantum dynamics:

$$\hat{H} = \sum_{j=0}^{d} \alpha_j \hat{H}_j = \alpha_0 \hat{I} + \alpha_1 \hat{\sigma}_z^{(0)} + \alpha_2 \hat{\sigma}_z^{(1)} + \alpha_3 \hat{\sigma}_z^{(0)}\hat{\sigma}_z^{(1)} + \alpha_4 \hat{\sigma}_y^{(0)}\hat{\sigma}_y^{(1)} + \alpha_5 \hat{\sigma}_x^{(0)}\hat{\sigma}_x^{(1)}, \tag{24}$$

N	ϵ	t	$\vec{\varphi} = (\varphi_1, \varphi_2, ..., \varphi_N)$ for $h(\theta) = t\sin(\theta)$
2	10^{-2}	0.0707	(-1.61, 1.67)
4	10^{-2}	0.311	(-1.03, 2.54, 1.36, -2.11)
8	10^{-2}	1.20	(-2.63, 1.66, -2.42, -1.79, 2.17, 1.43, -2.71),
16	10^{-2}	3.78	(0.23, 1.17, 3.07, -1.63, -1.78, 2.88, 1.71, 2.77, -1.95, 3.12, 1.97, -2.56, -2.87, 2.00, -2.29, 2.92, -0.51)
32	10^{-2}	10.1	(-2.94, 2.64, -2.32, 2.42, -2.86, -2.72, 2.4, -2.57, -2.96, 2.43, -2.53, -2.63, 2.33, 3.11, -2.17, 3.07, 2.24, -2.98, -1.99, -2.77, 2.22, 2.33, -2.62, -1.75, -2.15, 2.84, 1.68, 1.48, 2.46, -2.26, -0.92, -0.21)
2	10^{-4}	0.0070711	(-1.574, 1.5817)
4	10^{-4}	0.066948	(-0.6741, 2.5806, 0.7772, -2.4675)
8	10^{-4}	0.47498	(-0.0914, -1.2861, 2.1995, 0.995, -2.637, -1.5369, 1.9236, -3.0502)
16	10^{-4}	2.2164	(-0.5228, -2.9239, 1.3204, 2.6065, -1.73, -2.6191, 1.7225, 3.0121, -1.2286, -2.5276, 1.8058, -0.4605, 1.1563, -2.3492, 2.1623, -2.6188)
32	10^{-4}	7.3957	(1.3594, -2.7039, -1.5041, -1.0845, 2.4817, 2.5571, -3.0846, 1.0661, 2.6256, -2.0267, -2.0383, 3.0857, 1.9338, 2.6966, -2.1759, -2.562, 2.2839, 2.6493, -2.2281, -2.9435, 2.1726, -2.9934, -2.9327, 1.7167, -3.0853, -0.9383, -0.2802, -0.2376, -2.9556, 2.7875, -2.2875, 1.7822)

Table 3: Table of phases implementing target function $h(\theta) = t\sin(\theta)$ in quantum signal processing. Errors quoted refer to $\|\langle+|_b \hat{V}_{\vec{\varphi}}|+\rangle_b - \hat{W}_{\text{ideal}}\| \leq \epsilon$ in Thm. 6.

where $\hat{\sigma}_k^{(j)}$ denotes a Pauli matrix $\hat{\sigma}_k$ on the j^{th} qubit, the α_k are real coefficients that are a function of the bond length R, and $\hat{H}$ is in hartree units. One major goal of quantum chemistry is estimating energies at 'chemical accuracy'. This is $\epsilon = 1.6 \times 10^{-3}$ hartrees, at which point reaction rates at room temperature can be realistically predicted. In the following, we drop the identity term, as it only contributes a predictable global phase that can be subtracted from energy estimates, and all units are set to 1.

We now provide a comparison of gate counts required to simulate time evolution by Eq. 24 for time $t = \frac{1}{\epsilon} = 625$ with error ϵ to achieve chemical accuracy in quantum phase estimation. A rough estimate is provided for the linear-combination-of-unitaries approach in Cor. 8, and this is compared to best-case estimates of the number of exponentials required by first-order Trotterization of $e^{-i\hat{H}t}$ applied in [39], and upper bounds from higher-order Trotter-Suzuki formulas. Although the commonly used first-order Trotterization has a gate complexity that scales poorly like $\mathcal{O}(d^3t^2/\epsilon)$, compared to the $\mathcal{O}(d(\alpha t + \log(1/\epsilon)))$, where $\alpha = \sum_{j=1}^{d} |\alpha_j|$, of our approach, it has small constant factors in best case scenarios, uses no ancilla qubits, and establishes a baseline.

The first-order Trotterization of $e^{-i\hat{H}t}$ is

$$\left(e^{-i\hat{H}_1 t/n} \cdots e^{-i\hat{H}_d t/n}\right)^n = e^{-i\hat{H}t} + \mathcal{O}\left(\frac{(dt\|\hat{H}_{\vec{j}}\|)^2}{2n}\right), \quad \|\hat{H}_{\vec{j}}\| = \max\{\|\hat{H}_j\|,\ j \in [d]\}\}. \tag{25}$$

The required number of Trotter steps $n \approx (dt\|\hat{H}\|)^2/(2\epsilon)$ is estimated by solving $\epsilon = \mathcal{O}\left((dt\|\hat{H}_{\vec{j}}\|)^2/2n\right)$. This can be improved with higher-ordered Trotter-Suzuki product formulas [6] which bound the required number of exponentials $N_{\text{TS}} \leq 2d^2 t\|\hat{H}_{\vec{j}}\| e^{2\sqrt{\ln 5 \ln(dt\|\hat{H}_{\vec{j}}\|/\epsilon)}}$. We also compute the true error of Eq. 25 $\epsilon_{\text{T}} = \max_{|\psi\rangle} \|((e^{-i\hat{H}_1 t/n} \cdots e^{-i\hat{H}_d t/n})^n - e^{-i\hat{H}t})|\psi\rangle\|$, and find the minimum number of exponentials $N_{\text{T}} = dn$ such that $\epsilon_{\text{T}} = \epsilon$. This is the best-case estimate, as this computation is

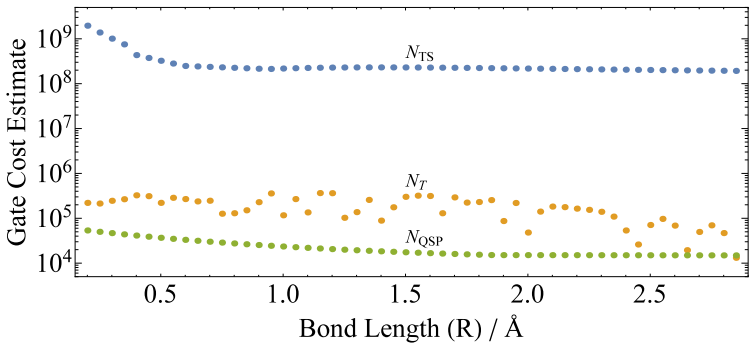

Figure 2: Estimated gate counts for achieving chemical accuracy in simulating the hydrogen dissociation Hamiltonian of Eq. 24. A comparison is made between upper bounds N_{TS} by ancilla-free Trotter-Suzuki methods (blue), a best-case first-order Trotterization N_{T} given inefficient classical pre-computation (yellow), and a reasonable estimate of our algorithm in the quantum signal processing formulation N_{QSP} (green).

in general intractable for larger systems.

Applying Cor. 8 requires the signal oracle $\hat{U} = \sum_{j=1}^{d} |j\rangle\langle j| \otimes \mathrm{sign}(\alpha_j)\hat{H}_j$. The signal state is $|G\rangle = \sum_{j=1}^{d} \sqrt{\frac{|\alpha_j|}{\alpha}}|j\rangle$. It is straightforward to verify that these satisfy the conditions $\langle G|\hat{U}|G\rangle = \hat{H}/\alpha$ and $\langle G|\hat{U}^2|G\rangle = \hat{I}$ of Thm. 2, thus $\hat{U}$ is already qubitized – this is a general feature of the Pauli basis, which is also unitary. We take the gate complexity of implementing $\hat{U}$ and the state oracle $\hat{G}$ to be $\approx d$ each. Thus the iterate $\hat{W}$ in Eq. 4 requires $\approx 3d$ gates. The number of iterates for quantum signal processing can be obtained from Fig. 1. In Eq. 24, $\alpha = \mathcal{O}(1)$. Thus $\alpha t = \alpha 625 \gg \log(1/\epsilon)$ is in the long time limit, so ≈ 2 iterates, or $6d$ gates per unit of simulation time is required, as $\hat{U}$ is aready qubitized. Thus the gate complexity of the Hamiltonian simulation with a linear combination of untiaries in algorithm is $N_{\mathrm{QSP}} \approx 6d\alpha t$.

These estimated gate costs for achieving chemical accuracy are plotted in Fig. 2 as a function of bond length R. The α_j parameters as a function of R are drawn from Table. 1 of [39]. While the estimate N_{QSP} is quite optimistic as it does not account for the constant factor costs of compiling multiply-controlled gates, even a factor of ~ 20 overhead keeps it competitive with the best-case Trotterization, which requires inefficient pre-computation. In any case, given minimal information about the structure of $\hat{H}$, the gate cost of our approach is several orders of magnitude better than ancilla-free Trotter-Suzuki methods at chemical accuracy, and will only improve further in simulations for longer times, at greater precision, and with more terms.

References

[1] Aaronson, S. The learnability of quantum states. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 463*, 2088 (2007), 3089–3114.

[2] Abramowitz, M., Stegun, I. A., et al. Handbook of mathematical functions. *Applied mathematics series 55* (1966), 62.

[3] Aharonov, D., and Ta-Shma, A. Adiabatic quantum state generation and statistical zero knowledge. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing* (New York, NY, USA, 2003), STOC '03, ACM, pp. 20–29.

[4] Babbush, R., Berry, D. W., Kivlichan, I. D., Wei, A. Y., Love, P. J., and Aspuru-Guzik, A. Exponentially more precise quantum simulation of fermions in second quantization. *New Journal of Physics 18*, 3 (2016), 033032.

[5] Barends, R., Kelly, J., Megrant, A., Veitia, A., Sank, D., Jeffrey, E., White, T. C., Mutus, J., Fowler, A. G., Campbell, B., Chen, Y., Chen, Z., Chiaro, B., Dunsworth, A., Neill, C., O'Malley, P., Roushan, P., Vainsencher, A., Wenner, J., Korotkov, A. N., Cleland, A. N., and Martinis, J. M. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature 508*, 7497 (Apr. 2014), 500–503.

[6] Berry, D. W., Ahokas, G., Cleve, R., and Sanders, B. C. Efficient quantum algorithms for simulating sparse hamiltonians. *Communications in Mathematical Physics 270*, 2 (2007), 359–371.

[7] Berry, D. W., and Childs, A. M. Black-box hamiltonian simulation and unitary implementation. *Quantum Info. Comput. 12*, 1-2 (Jan. 2012), 29–62.

[8] Berry, D. W., Childs, A. M., Cleve, R., Kothari, R., and Somma, R. D. Exponential improvement in precision for simulating sparse hamiltonians. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing* (New York, NY, USA, 2014), STOC '14, ACM, pp. 283–292.

[9] Berry, D. W., Childs, A. M., Cleve, R., Kothari, R., and Somma, R. D. Simulating hamiltonian dynamics with a truncated taylor series. *Phys. Rev. Lett. 114* (Mar 2015), 090502.

[10] Berry, D. W., Childs, A. M., and Kothari, R. Hamiltonian simulation with nearly optimal dependence on all parameters. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on* (Oct 2015), pp. 792–809.

[11] Berry, D. W., Cleve, R., and Gharibian, S. Gate-efficient discrete simulations of continuous-time quantum query algorithms. *Quantum Info. Comput. 14*, 1-2 (Jan. 2014), 1–30.

[12] Berry, D. W., and Novo, L. Corrected quantum walk for optimal hamiltonian simulation. *arXiv preprint arXiv:1606.03443* (2016).

[13] Blume-Kohout, R. Optimal, reliable estimation of quantum states. *New Journal of Physics 12*, 4 (2010), 043034.

[14] Boyd, J. P. Rootfinding for a transcendental equation without a first guess: Polynomialization of kepler's equation through chebyshev polynomial expansion of the sine. *Applied Numerical Mathematics 57*, 1 (2007), 12 – 18.

[15] Brandao, F. G., and Svore, K. Quantum speed-ups for semidefinite programming. *arXiv preprint arXiv:1609.05537* (2016).

[16] Childs, A. M. On the relationship between continuous- and discrete-time quantum walk. *Commun. Math. Phys. 294*, 2 (2010), 581–603.

[17] Childs, A. M., Cleve, R., Deotto, E., Farhi, E., Gutmann, S., and Spielman, D. A. Exponential algorithmic speedup by a quantum walk. In *Proceedings of the Thirty-fifth Annual*

ACM Symposium on Theory of Computing (New York, NY, USA, 2003), STOC '03, ACM, pp. 59–68.

[18] Childs, A. M., and Kothari, R. Limitations on the simulation of non-sparse hamiltonians. *Quantum Info. Comput. 10*, 7 (July 2010), 669–684.

[19] Childs, A. M., Kothari, R., and Somma, R. D. Quantum linear systems algorithm with exponentially improved dependence on precision. *arXiv preprint arXiv:1511.02306* (2015).

[20] Childs, A. M., and Wiebe, N. Hamiltonian simulation using linear combinations of unitary operations. *Quantum Info. Comput. 12*, 11-12 (Nov. 2012), 901–924.

[21] Chowdhury, A. N., and Somma, R. D. Quantum algorithms for gibbs sampling and hitting-time estimation. *arXiv preprint arXiv:1603.02940* (2016).

[22] Daskin, A., and Kais, S. An ancilla based quantum simulation framework for non-unitary matrices. *arXiv preprint arXiv:1606.04315* (2016).

[23] Debnath, S., Linke, N. M., Figgatt, C., Landsman, K. A., Wright, K., and Monroe, C. Demonstration of a small programmable quantum computer with atomic qubits. *Nature 536*, 7614 (Aug. 2016), 63–66.

[24] Flammia, S. T., Gross, D., Liu, Y.-K., and Eisert, J. Quantum tomography via compressed sensing: error bounds, sample complexity and efficient estimators. *New Journal of Physics 14*, 9 (2012), 095022.

[25] Grover, L. K. A fast quantum mechanical algorithm for database search. 212–219.

[26] Harrow, A. W., Hassidim, A., and Lloyd, S. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett. 103* (Oct 2009), 150502.

[27] Karam, L. J., and McClellan, J. H. Chebyshev digital fir filter design. *Signal Process. 76*, 1 (1999), 17 – 36.

[28] Kimmel, S., Lin, C. Y.-Y., Low, G. H., Ozols, M., and Yoder, T. J. Hamiltonian simulation with optimal sample complexity. *arXiv preprint arXiv:1608.00281* (2016).

[29] Kivlichan, I. D., Wiebe, N., Babbush, R., and Aspuru-Guzik, A. Bounding the costs of quantum simulation of many-body physics in real space. *arXiv preprint arXiv:1608.05696* (2016).

[30] Kothari, R. *Efficient algorithms in quantum query complexity.* PhD thesis, 2014.

[31] Lloyd, S. Universal quantum simulators. *Science 273*, 5278 (Aug 23 1996), 1073.

[32] Lloyd, S., Mohseni, M., and Rebentrost, P. Quantum principal component analysis. *Nat Phys 10*, 9 (Sept. 2014), 631–633.

[33] Low, G. H., and Chuang, I. L. Optimal hamiltonian simulation by quantum signal processing. *arXiv 1606* (Jun 2016), 02685.

[34] Low, G. H., Yoder, T. J., and Chuang, I. L. Quantum imaging by coherent enhancement. *Phys. Rev. Lett. 114* (Mar 2015), 100801.

[35] Low, G. H., Yoder, T. J., and Chuang, I. L. The methodology of composite quantum gates. *arXiv 1603* (Mar 2016), 03996.

[36] McClellan, J., Parks, T., and Rabiner, L. A computer program for designing optimum FIR linear phase digital filters. *IEEE Trans. Audio Electroacoust. 21*, 6 (Dec 1973), 506–526.

[37] Meinardus, G. *Approximation of functions: Theory and numerical methods*, vol. 13. Springer Science & Business Media, 2012.

[38] O'Donnell, R., and Wright, J. Efficient quantum tomography. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing* (New York, NY, USA, 2016), STOC 2016, ACM, pp. 899–912.

[39] O'Malley, P. J. J., Babbush, R., Kivlichan, I. D., Romero, J., McClean, J. R., Barends, R., Kelly, J., Roushan, P., Tranter, A., Ding, N., Campbell, B., Chen, Y., Chen, Z., Chiaro, B., Dunsworth, A., Fowler, A. G., Jeffrey, E., Lucero, E., Megrant, A., Mutus, J. Y., Neeley, M., Neill, C., Quintana, C., Sank, D., Vainsencher, A., Wenner, J., White, T. C., Coveney, P. V., Love, P. J., Neven, H., Aspuru-Guzik, A., and Martinis, J. M. Scalable quantum simulation of molecular energies. *Phys. Rev. X 6* (Jul 2016), 031007.

[40] Poulin, D., Hastings, M. B., Wecker, D., Wiebe, N., Doberty, A. C., and Troyer, M. The trotter step size required for accurate quantum simulation of quantum chemistry. *Quantum Info. Comput. 15*, 5-6 (Apr. 2015), 361–384.

[41] Reiher, M., Wiebe, N., Svore, K. M., Wecker, D., and Troyer, M. Elucidating reaction mechanisms on quantum computers. *arXiv preprint arXiv:1605.03590* (2016).

[42] Somma, R. D. A trotter-suzuki approximation for lie groups with applications to hamiltonian simulation. *Journal of Mathematical Physics 57*, 6 (2016).

[43] Szegedy, M. Spectra of quantized walks and a $\sqrt{\delta\epsilon}$ rule. *arXiv preprint quant-ph/0401053* (2004).

[44] Wecker, D., Bauer, B., Clark, B. K., Hastings, M. B., and Troyer, M. Gate-count estimates for performing quantum chemistry on small quantum computers. *Phys. Rev. A 90* (Aug 2014), 022305.

[45] Yoder, T. J., Low, G. H., and Chuang, I. L. Fixed-point quantum search with an optimal number of queries. *Phys. Rev. Lett. 113* (Nov 2014), 210501.

[46] Yung, M.-H., Whitfield, J. D., Boixo, S., Tempel, D. G., and Aspuru-Guzik, A. *Introduction to Quantum Algorithms for Physics and Chemistry.* John Wiley & Sons, Inc., 2014, pp. 67–106.